REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
OCT 08	Conference Paper Postprint	Oct 06 - Oct 07

4. TITLE AND SUBTITLE
SEMANTIC INTEROPERABILITY IN DISTRIBUTED PLANNING

5a. CONTRACT NUMBER In-House
5b. GRANT NUMBER N/A
5c. PROGRAM ELEMENT NUMBER 62702F

6. AUTHOR(S)
Gennady R. Staskevich, Jeffrey W. Hudack, Joseph A. Carozzoni and James H. Lawton

5d. PROJECT NUMBER 558S
5e. TASK NUMBER IH
5f. WORK UNIT NUMBER DP

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
AFRL/RISB
525 Brooks Rd.
Rome, NY 13441-4505

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AFRL/RISB
525 Brooks Rd.
Rome, NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
N/A

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-RI-RS-TP-2008-4

12. DISTRIBUTION AVAILABILITY STATEMENT
Approved for public release: distribution is unlimited. PA# 08-0088

13. SUPPLEMENTARY NOTES

Presented at the International Command and Control Research and Technology Symposium 2008, Bellevue, WA Jun 17-19,2008

14. ABSTRACT
The USAF Command and Control (C2) is undergoing a transformation from a co-located, theater-centric process to one that is worldwide and distributed. A key challenge for this transformation to Globally-Linked Air and Space Operations Centers is developing the ability to collaboratively plan and execute operation with multiple cooperating command centers. This paper describes an in-house program underway at the USAF Research Laboratory Information Directorate that is developing technologies to support the concepts of Network Centric Operations. In particular, research is presented that extends the Object Model Working Group's Core Plan Representation (CPR) framework utilizing semantic technologies to capture planning experiences in both human- and machine-readable form. A key feature of these extensions is common,interoperable plan representation amongst the distributed heterogeneous planning agents. Semantic interoperability of the plan representation is critical to support distributed planning. The initial approach to achieving interoperability is a limited taxonomy for describing key plan-related information. The research presented utilizes open standards semantic technology to encapsulate plans as self-describing semantic objects.

15. SUBJECT TERMS
Distributed Command and Control, Analogical Reasoning, Network Centric Planning, Intelligent Agents

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
U	U	U	UU	22	Joseph A. Carozzoni
					19b. TELEPHONE NUMBER (*Include area code*) N/A

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Title: "Semantic Interoperability in Distributed Planning"

Track 1: C2 Concepts, Theory, and Policy

Authors

Gennady R Staskevich

Jeffrey W. Hudack

Dr. James Lawton

Joseph A. Carozzoni

Point of Contact:

Gennady R Staskevich

Air Force Research Laboratory Information Directorate

Systems & Information Interoperability Branch (RISE)

525 Brooks Road, Rome, NY 13441-4505

315-330-4889

Gennady.Staskevich@rl.af.mil

Postprint

Semantic Interoperability in Distributed Planning

Abstract

The USAF Command and Control (C2) is undergoing a transformation from a co-located, theater-centric process to one that is worldwide and distributed. A key challenge for this transformation to Globally-Linked Air and Space Operations Centers is developing the ability to collaboratively plan and execute operations with multiple cooperating command centers. This paper describes an in-house program underway at the USAF Research Laboratory Information Directorate that is developing technologies to support the concepts of Network Centric Operations. In particular, research is presented that extends the Object Model Working Group's *Core Plan Representation* (CPR) framework utilizing semantic technologies to capture planning experiences in both human- and machine-readable form. A key feature of these extensions is common, interoperable plan representation amongst the distributed heterogeneous planning agents. Semantic interoperability of the plan representation is critical to support distributed planning. The initial approach to achieving interoperability is a limited taxonomy for describing key plan-related information. The research presented utilizes open standards semantic technology to encapsulate plans as self-describing semantic objects.

Keywords: Distributed Planning, Semantic Interoperability, Core Plan Representation, Network-Centric Operations (NCO), Mix-Initiative Planning

Abstract 2
Introduction 4
1.1 Future C2 Requirements 4
1.2 Objective of the DEEP Project 5
1.3 Semantic Interoperability as a key research topic 5
2 Framework for Supporting Distributed, Mixed-Initiative Planning 6
2.1 DEEP Overview 6
2.2 The Original ARPI CPR 7
2.3 DEEP adaptation of CPR 8
3 Leveraging Semantic Technologies 13
3.1 DEEP-CPR Semantic Extensions 13
3.1.1 URI Reference 14
3.1.2 RDF Metadata Model 15
3.2 Semantic Technology Benefits and Challenges 16
4 Conclusion 18
References 20

Introduction

This paper describes an in-house program underway at the U.S. Air Force Research Laboratory Information Directorate known as Distributed Episodic Exploratory Planning (DEEP). DEEP is developing technologies to support the concepts of *Network Centric Operations* (NCO) [Alberts & Garstka & Stein, 1999]. In particular, this paper presents research into extending the Object Model Working Group's *Core Plan Representation* (CPR) framework [Pease 1998] to capture planning experiences in both human- and machine-readable form, as well as to provide semantic interoperability among distributed, heterogeneous planner. We begin with a discussion of the way Command and Control (C2) is changing, which motivates the need for a system like DEEP. Within this paper, we use the term *semantic* to describe an agreed upon, controlled vocabulary that ensures a shared understanding of concepts and allows for unambiguous communication between agents, both human and machine within a distributed environment.

1.1 Future C2 Requirements

To meet future challenges, U.S. forces are in the midst of a "transformation" to not only support traditional high-tempo, large force-on-force engagements, but also smaller-scale conflicts characterized by insurgency tactics and time-sensitive targets of opportunity. This transformation requires a vastly new Command and Control (C2) process that can adapt to the any level of conflict, provides a full-spectrum joint warfighting capability, and can rapidly handle any level of complexity and uncertainty. In addition to the air and space domains, the U.S. Air Force is standing up to the challenges of the cyber domain. A key challenge of cyber C2 is the speed at which electrons move, requiring a C2 system of unprecedented response time, global arena, and human expertise that may not be located in a single command center.

To meet these future challenges, the U.S. Air Force (USAF) is moving towards a model of continuous air operations not bounded by the traditional 24-hour Air Tasking Order (ATO) cycle. Meeting these objectives will require a highly synchronized, distributed planning and replanning capability. Experience with recent operations also reveals that the C2 process must transition from a process of observation and reaction to one of prediction and preemption. As a potential way ahead, AF/A5 (Plans) in May 2006 released a revolutionary vision paper titled "C2 Enabling Concepts" depicting what a potential future C2 environment could be. Four key concepts emerged from this vision of a future AOC:

- Distributed/Reachback planning
- Redundant/Backup planning
- Continuous planning
- Flexible, scalable, tailorable C2

The research presented in this paper has been focused on two emerging concepts for the future of C2. First, developing a C2 environment that supports the vision of Network Centric Operations (NCO). The tenets of NCO are:

- Information sharing
- Shared situational awareness
- Knowledge of commander's intent

Second, developing a distributed C2 environment that supports Cyber Warfare. We currently assume that the cyber domain requires a faster than a real-time (i.e., predictive) C2 capability that is not bounded by traditional thinking (i.e., air and space). The implications being an all encompassing, global battlespace that requires expertise that is seldom co-located.

1.2 Objective of the DEEP Project

The long-term goal of the Distributed Episodic Exploratory Planning (DEEP) project is to develop in-house a prototype system for distributed, mixed-initiative planning that improves decision-making by applying analogical reasoning over an experience base. The two key objectives of DEEP are:

- Provide a mixed-initiative planning environment where human expertise is captured and developed, then adapted and provided by a machine to augment human intuition and creativity.
- Support the distributed planners in multiple cooperating command centers to conduct distributed and collaborative planning.

The architecture of DEEP was explicitly designed to support the tenets of NCO in a true distributed manner. Because DEEP is not based on any current C2 system, we are able to explore concepts such as combining planning and execution to support dynamic replanning, machine-mediated self synchronization of distributed planners, and experiment with the impact of trust in an NCO environment (i.e., "Good ideas are more important than their source").

1.3 Semantic Interoperability as a key research topic

Alberts and Hayes (2007) advocate bold new approaches beyond current organizational process, focusing on what is possible for NCO. Their suggested recommendation is to systematically explore the following high priority research topic areas:

1. Taxonomy for planning and plans;
2. Quality metrics for planning and plans;
3. Factors that influence planning quality;
4. Factors that influence plan quality;
5. Impact of planning and plan quality on operations;
6. Methods and tools for planning; and
7. Plan visualization

In order to achieve the vision of DEEP, essentially all the above topics need to be addressed. The first topic was the starting point and has received the majority of our attention. The earliest effort in support of distributed planning was on the CPR, an object-oriented plan framework developed under the ARPA-Rome Laboratory Planning Initiative (ARPI). CPR is based on the Unified Modeling Language (UML), which is well suited as the human-machine dialog to support mixed-initiative planning. The recursive nature of CPR supports multi-level planning at all levels (strategic, operational, and tactical), along with plan fragments supporting distributed planning on a plan simultaneously. A key research topic for DEEP, addressed by the work presented in this paper, is maintaining referential integrity when distributed planners simultaneously work on multiple sub-plans and/or plan fragments of a larger plan.

2 Framework for Supporting Distributed, Mixed-Initiative Planning

In this section we discuss our initial approach to supporting distributed, mixed-initiative planning for C2 activities using NCO principles. We first present an overview of the DEEP architecture along with a description of the CPR framework as developed in ARPI. We then describe the extensions to CPR that were needed to support the goals of the DEEP project. In Section 3 we discuss our plans to further enhance CPR using state-of-the-art semantic technology.

2.1 DEEP Overview

As shown in Figure 1, DEEP is a systems-of-systems architecture, comprised of the following systems:

- Distributed Blackboard for multi-agent, non-deterministic, opportunistic reasoning
- Case-Based Reasoning system to capture experiences (successes and/or failures)
- Episodic Memory for powerful analogical reasoning
- Multi-Agent System for mixed-initiative planning
- ARPI CPR for human-to-machine common dialog
- Constructive Simulation for exploration of plausible future states

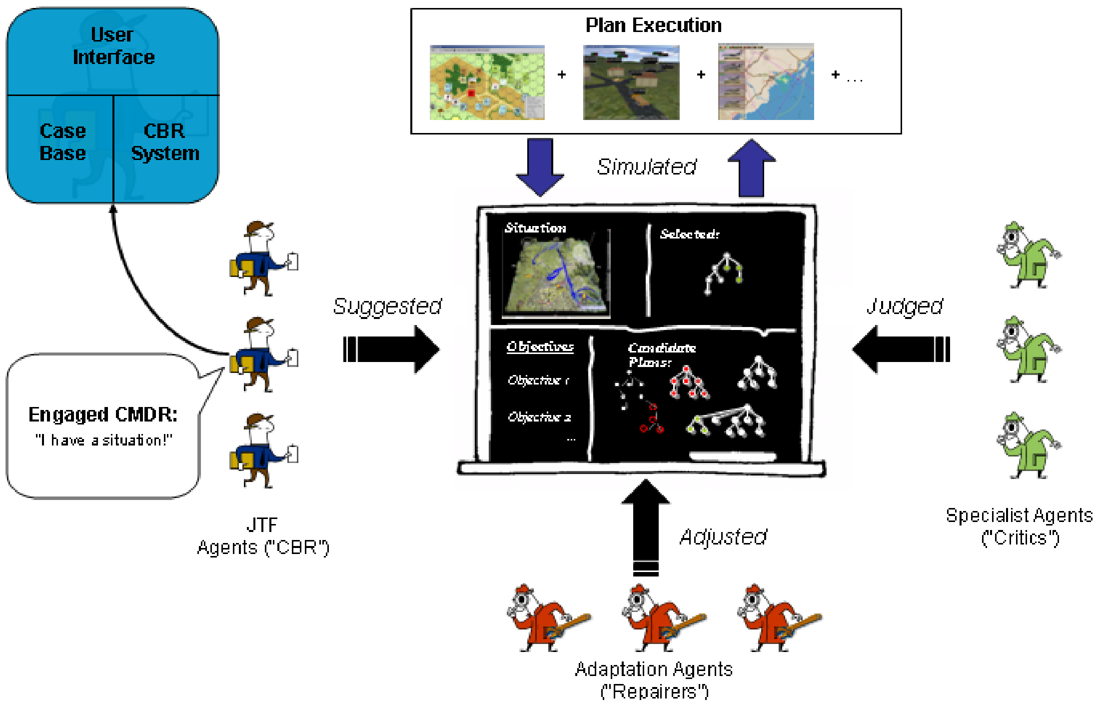

Figure 1 - DEEP Architecture

A prime motivation for the DEEP architecture was maximizing collaboration in planning. In the information age, the term "collaboration" has taken on many roles. A common type of collaboration is the chat-room where multiple parties exchange text messages in an asynchronous mode. This type of collaboration was considered insufficient for information age warfare and the term distributed planning substituted. Our definition of distributed planning includes both asynchronous and synchronous collaboration. An initial challenge was prohibiting the distributed planning aspect of DEEP to degrade to asynchronous chat-room text messaging. In defending against this, for development purposes only, an artificial barrier to human-to-human collaboration was imposed, forcing all interaction to be machine-mediated. Additionally, this requirement will facilitate the mixed-initiative aspect of planning.

2.2 The Original ARPI CPR

The Advanced Research Projects Agency (ARPA, now known as DARPA) and the Air Force's Rome Laboratory (RL, now part of the Air Force Research Laboratory) conducted extensive research in plan representation under the ARPA-RL Planning Initiative (ARPI). One of the more promising results of that research was the CPR, an object-oriented plan framework based on the Unified Modeling Language (UML). The prime motivation behind CPR was plan interoperability. CPR offered three critical capabilities to support distributed planning. First, being UML-based, it is well suited as the human-machine dialog to support mixed-initiative planning. Second, the recursive nature of CPR supports planning at all levels (strategic, operational, and tactical), along with the inclusion of plan fragments that can support distributed

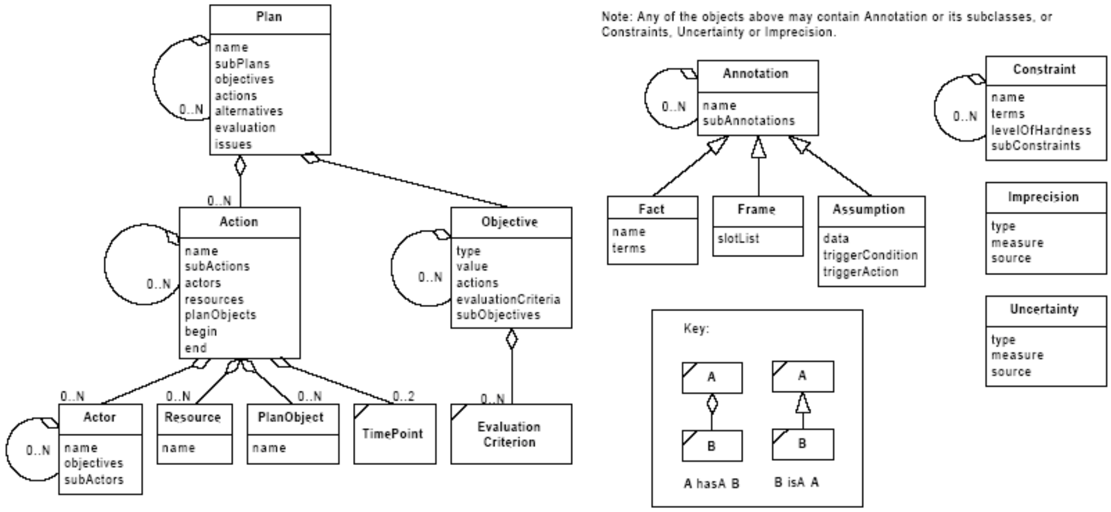

Figure 2 - ARPI-CPR Model

simultaneous planning. Lastly, plan interoperability leads to the ability to support full-spectrum C2 (air, space, and cyber) as well as multi-service (i.e., joint) C2.

The Air Force Research Laboratory co-sponsored the Object Modeling Group (OMG) to derive a "basic level" plan representation for the domain of planning in order to facilitate information exchange among different planning systems. Their study resulted in what is known as the ARPI-CPR model [Pease & Carrico 1997].

The ARPI-CPR model, shown in Figure 2, is an abstract specification that provides a highly flexible and recursive architecture for the plan representations. The object-oriented design of CPR is based on a commonly shared set of objects and intended to be extensible: Action, Actor, Objective and Resource [Pease & Carrico 1997]. This framework syntactically captures the foundational planning concepts using object-oriented design.

2.3 DEEP adaptation of CPR

Although the DEEP framework has adopted ARPI-CPR model as a basis for its plan representation, the model specification was too abstract to be used directly. Our extensions to the APRI-CPR model were driven by the need to use a richer content description for reasoning about plans, as well as to be able to encode plans as part of cases in a CBR system. A significant difference between ARPI- and DEEP-CPR is that planning information within DEEP is structured (currently using taxonomies), making the free text used in ARPI-CPR inadequate. Further, since DEEP uses a CBR system for plan selection and storage, it was necessary to extend DEEP-CPR to make it a component of a case.

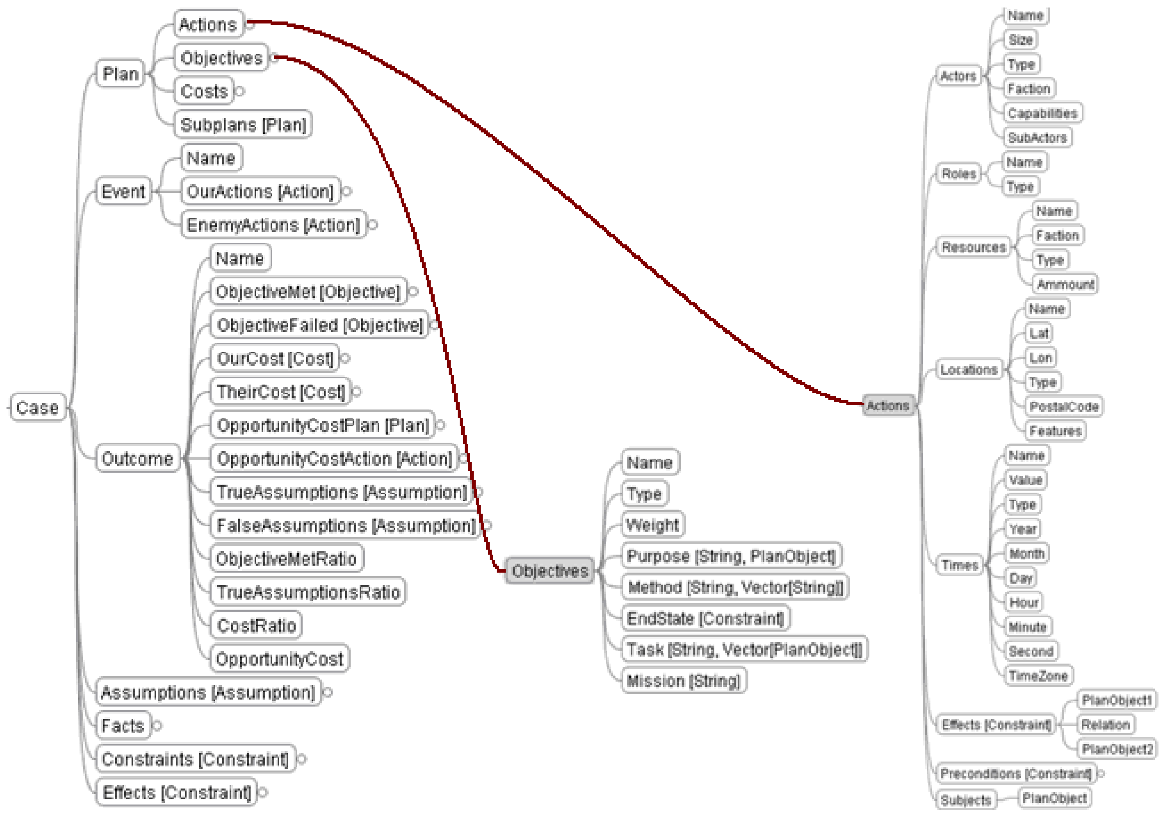

Figure 3 - DEEP-CR Model

Figure 3 presents a representation of the CPR model as extended for DEEP. At the most abstract level planning experiences are encoded in DEEP at *Cases* in a CBR system, which we refer to as the DEEP Case Representation (DEEP-CR).

- ***Case*** **–** An encapsulation of a planning experience (DEEP-CR). The highest level of abstraction that captures the plan *and* the situational awareness surrounding the plan. It is an integral part of the case base as required by the CBR system.

A case is composed of as three components: a *Plan*, an *Event*, and an *Outcome*. The plan is represented in DEEP-CPR. Events and outcomes provide contextual information about the experience of executing the associated plan.

- ***Plan*** **–** A concept that is defined by a set of objectives and the corresponding actions that address them while including the situational constraints such as costs, and other sub plans. The plan is applicable at any level of command chain whether it is strategic, operational, or tactical.

- ***Event*** **–** A concept that captures significant action occurrences into 2 categories: local (e.g., “our”) actions and enemy actions. The purpose of including events is to enable the capability to capture the cause and effect of the associated plan.

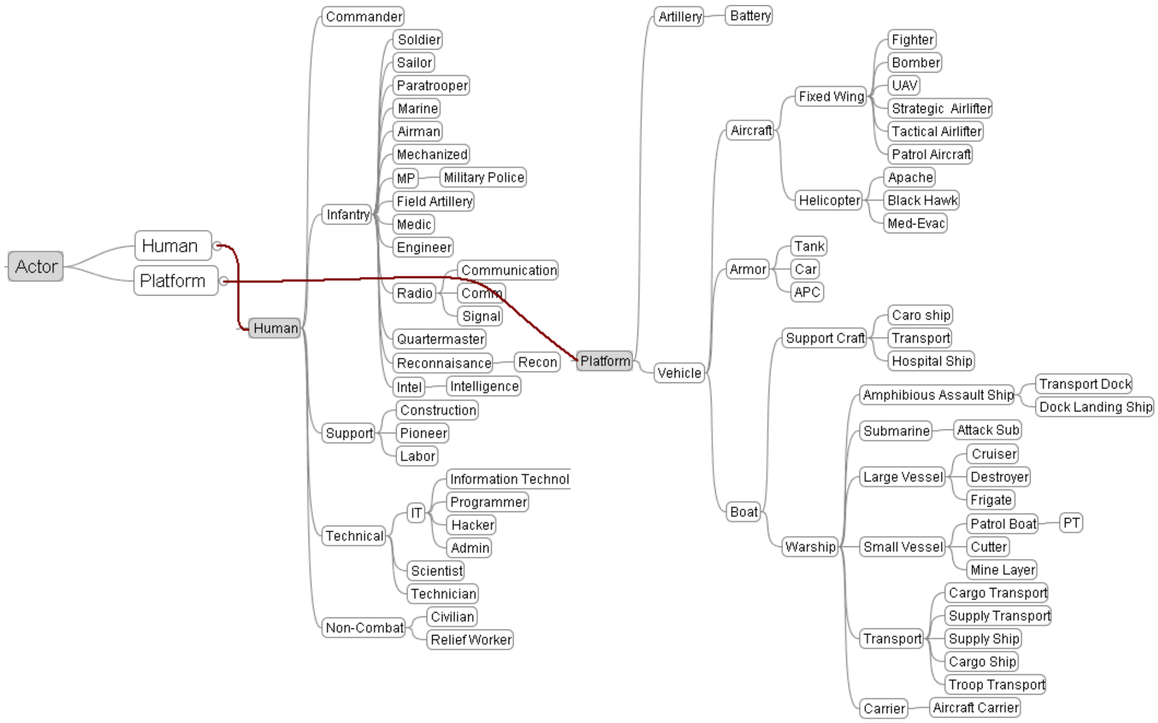

Figure 4 - Actor Taxonomy

- ***Outcome*** – A concept that captures actual events, costs and other supporting information regarding the results of executing the associated plan. This concept enables a subjective evaluation of the overall effectiveness or ineffectiveness of candidate plan. An outcome may include such information as costs (local and enemy) and assumptions of the associate plan.

As mentioned above, a Plan within the DEEP-CR structure is where the actual plan is stored, and is represented using our extended version of CPR, DEEP-CPR. Just as in the original ARPI-CPR model, the main components of a plan represented with DEEP-CPR are *Actions*, *Actors*, *Resources* and *Objectives*.

- ***Action*** **–** A concept that describes an activity being performed by an Actor. It is an integral part of a plan representation as identified by the ARPI-CPR, and plays an important role during the retrieval and revision stages of the CBR cycle.

- ***Actor*** **–** An Actor is an entity that carries out an associated action. An Actor in one Action may also be a Resource in another (e.g., infantry). Actors are taxonomically described (see Figure 4).

- ***Resource*** **–** Similar to Actor, a concept that is used to support the performance the Action by an Actor. Similar to Actors, Resources are described using taxonomies (see Figure 5).

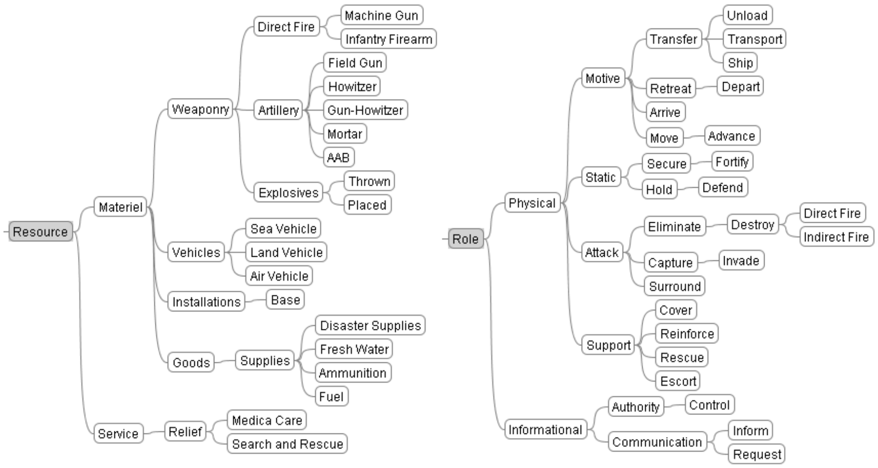

Figure 5 – Resource and Role Taxonomies

- ***Objective –*** A concept that defines the primary purpose or a motivation behind performing one or more Action to achieve some goal.

While the ARPI-CPR model identifies the need for these key concepts, it does not specify how they should be represented. As such, DEEP-CPR includes a number of extensions to these components to support our planning environment. For many of these extensions, we have also defined corresponding taxonomies of values that would fill the associated extended plan structures. Example taxonomies for Actors, Resources and Roles are shown in Figure 4 and Figure 5. The use of taxonomies was needed to support our analogical reasoning mechanism, by which the DEEP system is able to select the most appropriate past experiences to apply to new problems. In Section 3, we present our approach to using *semantic technologies* to address this problem.

In addition to Actors and Resources (described above) the extensions for plan Actions include: *Roles, Subjects, Locations, Times, Effects, and Preconditions*:

- ***Role –*** A concept that defines a proper or a customary function of an action. The purpose of a *Role* is to store taxonomically defined verbs to succinctly describe a function of an action (see Figure 5).

- ***Subject –*** The intended target of an action, essentially any CPR object.

- ***Locations and Times –*** While ARPI-CPR used *TimePoint* for these concepts, within our DEEP-CPR adaptation, they are split into the individual components *Location* and *TimeSpec*:

    - A ***Location*** is a concept that captures both the physical coordinate of a given location along with certain defining characteristics by using taxonomically defined features such as: land, sea, shore, forest, hill, cave

etc. Within a location, *features* are used to enable a higher-precision case matching and higher-fidelity adaptation.

  - A ***TimeSpec*** is a concept that is used for capturing and describing the temporal information within DEEP. It is composed of the values: Year, Month, Day, Hour, Minute, Second and Time zone. These values concisely capture the fundamental aspects of temporal information.

- ***Effects, and Preconditions –*** Both concepts are a type of constraint. A constraint is essentially a triplet composed of two objects and a relationship between them [Brickley & Guha, 2002]. These concepts are used to enable a higher-fidelity description of a plan and its surrounding information within DEEP.

  - ***Effect –*** A constraint on the world that will result as a consequence of an action.

  - ***Precondition –*** A constraint that needs to exist in the world prior to performing of an action.

The DEEP-CPR extensions to the Objective concept are based on a military doctrine developed by the US Army and Marine Corps. They are composed of three main parts: *Purpose, Method* and *End State* [FM101-5 US Army Field Manual, 1997]. The Purpose states what you are trying to achieve. The Method states how are you going to achieve the purpose. The End State defines a measurable metric to determine the success or a failure of a Purpose. Further, we have similarly expanded Method based on the military doctrine [LtGen M. R. Berndt (Ret.), 2007] to include two more levels of abstraction, *Mission* and *Task,* enabling the commanding officer to state a required level of specificity on how a Method is to be achieved.

The Objective is expressed as a composition of multiple DEEP-CPR objects including: *Purpose, Method, Mission, Task,* and *End State*.

- ***Purpose –*** A concept that states what the plan is trying to achieve. A purpose is represented by its *Type*, which is taxonomically defined, and a *Subject*. A Subject may be any valid DEEP-CPR object.

- ***Method –*** A concept that states how a given Purpose is to be achieved. A Method is composed of a *Type* and a *Center of Gravity* (COG), both of which are taxonomically defined.

- ***Mission –*** A concept that classifies the type of mission the Plan accomplishes, drawn from a taxonomy of mission types.

- ***Task –*** A concept that specifies an assignment to perform a definitive piece of work. A task is composed of a taxonomically defined *Type* along with a *Subject*.

Figure 6 – Objective Decomposition

- ***End State –*** A concept that defines a measurable metric to determine a success or a failure of a given Purpose. An End State is described as a relationship (drawn from a taxonomy of relationships) on two Resources.

For example, in *Operation Unified Assistance*, the U.S. Military jointly assisted in the aftermath of the 2004 tsunami that devastated the Aceh region of Indonesia. One of the objectives was to *provide search and rescue support in the Aceh region*. This objective is broken down into *Purpose*, *Mission* and *End-State* statements as follows: *Purpose* - In order to Support Disaster Victims, *Mission* - Provide Aero-medical Evacuation, and *End-State* - Search and Rescue must be AT Aceh Province. Each statement is taxonomically described (see Figure 6).

3 Leveraging Semantic Technologies

As discussed in Section 2.3, the current version of DEEP-CPR uses a collection of informal taxonomies to represent plan-related information. While the terms used in these taxonomies have meaning to the people that developed them, they are in fact merely collections of symbols that have no explicit meaning to the machine. As such, the interpretation of this information must be programmed into the agents that interpret and manipulate plans represented in DEEP-CPR.

In this section we present an approach to extending DEEP-CPR to leverage semantic technologies. The goal of these extensions will be to allow plans in DEEP to be semantically self-describing. This will allow developers of DEEP components, such as plan analysis agents, to create domain-independent approaches free of hard-coded knowledge needed to interpret plans.

3.1 DEEP-CPR Semantic Extensions

As a first step, we intend to use RDF (Resource Description Framework) [Brickley & Guha, 2002] as the foundational layer for the semantic extensions to DEEP-CPR. RDF was selected for its simplicity and flexibility to capture and express of the semantics within a DEEP planning system. Furthermore, we will attempt to ground the foundational DEEP-CPR concepts with a commonly accepted upper ontology model such as Suggested Merged Upper Ontology (SUMO). [Pease, Niles & Li, 2002]

For example, the Purpose slot of a plan object would currently be filled in with a symbol from a taxonomy, such as the one shown in Figure 5. Using RDF, it might be represented as shown in Figure 7, which makes statements declaring that Objective is a subclass of, or part of, a Plan and Purpose is a subclass of an Objective. These statements do not define the meaning of these terms, but rather provide relationships between the terms within a controlled vocabulary.

```
<rdfs:Class rdf:about=”&cpr;Objective”
   rdfs:label=”Objective”>
   <rdfs:subClassOf rdf:resource=”&cpr;Plan”/>
</rdfs:Class>
<rdfs:Class rdf:about=”&cpr;Purpose”
   rdfs:label=”Purpose”>
   <rdfs:subClassOf rdf:resource=”&cpr;Objective”/>
</rdfs:Class>
```

Figure 7 - RDF Example

Given the difficulty of reading and interpreting RDF statements in XML, these triples can instead be represented in a more simple form:

```
<cpr:Objective  rdfs:subClassOf cpr:Plan/>
<cpr:Purpose    rdfs:subClassOf cpr:Objective />
```

3.1.1 URI Reference

To support these extensions, the text-based entries of the DEEP-CPR objects will first be converted to Uniform Resource Identifiers (URIs). A URI is the fundamental building block of semantic technologies. A URI is a character string that encodes a networked resource, providing a structured and stable method for representing concepts and resources in a distributed information space.

Unlike text-based identifiers, which are interpreted by each local system independently, a URI-based infrastructure provides a self-describing data model that ensures agreement on the concepts being used. This is similar to the pass-by-reference model used by modern-programming languages such as Java in which information is passed as a pointer to a value rather than a copy of the value. Therefore, when the concept is changed at the source it is not necessary to copy those changes to all instances since the pointer to that information does not need to change.

HTTP is the most common protocol used to create URIs as it is well-suited to resource indexing. In DEEP-CPR, a concept such as the Purpose of a plan could be represented as http://deep.af.mil/namespace/concepts#Purpose. The server containing the reference is *deep.af.mil*, with *concepts* being located in a namespace virtual directory on that server.

The *#Purpose* is interpreted as the reference to the Purpose concept within the concepts page.

3.1.2 RDF Metadata Model

Once the DEEP-CPR entries are converted to URIs, they will be encoded into RDF statements. RDF is an important component of semantic technologies being supported by the World Wide Web Consortium (W3C) and the semantic web community. It is a general-purpose language defined by the use of triples to identify and describe networked resources that are represented by URIs. An RDF statement is a triple (S, P, O) in which:

- S is a URI, the subject of the statement
- P is a URI, the predicate of the statement denoting a relationship
- O is either a URI or a literal (plain text) and represents the object or value of the predicate P for subject S

RDF Schema (RDF-S) [Brickley & Guha, 2003] is a vocabulary that introduces additional structure to RDF allowing domain-independent definition of classes, subclasses, and properties. These structural concepts provide a means for declaring relationships between concepts that are generally understood without specific knowledge of the domain. It is this fundamental understanding that allows shared understanding between information spaces without explicit mapping and transformation.

For example, assume the existence of the following namespaces for *rdfs* and *cpr*, that define the RDF-S and CPR vocabularies. We can then define a case using the abbreviated syntax for a Case called *Operation Unified Assistance*. As in the example shown in Figure 6, search and rescue assistance will be provided to victims of a disaster in the Aceh region. Concepts that have a numbered suffix, such as `Case2398`, represent the actual internal concept name for unique identification within the system. Since it is possible (even likely, given variations on the same plan) that two cases could have the same name we cannot rely on the name of the case to distinguish it from another case. Instead, the name becomes a property of the concept via the *hasName* relationship.

```
[Instance object types]
<Case2398 rdfs:type cpr:Case/>
<Obj8723 rdfs:type cpr:Objective/>
<Pur4423 rdfs:type cpr:Purpose/>
<Role6415 rdfs:type cpr:Role/>
<Entity0123875 rdfs:type cpr:People/>
<Loc23498 rdfs:type cpr:Location/>

[Object properties]
<Case2398 cpr:hasName "Operation Unified Assistance"/>
<Case2398 cpr:hasObjective Obj8723/>
<Obj8723 cpr:hasName "Provide Search and Rescue in Aceh Region"/>
<Obj8723 cpr:hasPurpose Pur4423/>
<Pur4423 cpr:hasName "Support Victims"/>
<Pur4423 cpr:hasRole Role6415/>
<Pur4423 cpr:hasSubject Entity0123875/>
<Role6415 cpr:hasType cpr:Support/>
```

```
<Entity0123875 cpr:hasName "Disaster victims in Aceh region"/>
<Entity0123875 cpr:hasLocation Loc23498/>
<Loc23498 cpr:hasName "Aceh">
```

Each of the concepts used will map to URIs that provide rich textual description defining the meaning of the term. This more explicit definition provides some guarantee that terms used will be based on agreed-upon meanings rather than local interpretation of a symbol. In addition, this self-describing format provides a set of constraints on parts of the plan that will support both comparison and reasoning for the system.

Unlike URIs, RDF statements cannot be stored on web servers as HTML pages. Therefore, it will be necessary to utilize RDF Stores within the DEEP architecture to provide storage of triples. One likely option will be to use an Oracle database, as Oracle will be including support for storage and query of the RDF format [Alexander et al., 2004] in their next major version. Unlike URIs, RDF statements are more challenging to access in a distributed environment, [Nejdl, Siberski & Sintek, 2003] although methods such as peer-to-peer [Cai & Frank, 2004] indexing have been investigated.

3.2 Semantic Technology Benefits and Challenges

The use of semantic technologies as data structures provides significant benefit to both computational and user-based interaction with the information being stored. These benefits include:

- *Expressive.* Since each piece of information is attached to a base hierarchy of structural and functional relationships it can be given more description than a simple string name can provide. Issues such as individual meaning and context can begin to be represented and attached to information to be sure that it is used correctly by the systems that will process it. However, it should be noted the commitment to some first-order logics such as those provided by OWL may also limit the expressivity that can be provided. [Pan & Horrocks, 2004]

- *Abstraction*. The hierarchical organization of formal ontologies provides a structure that can be directly interpreted as layers of abstraction. This provides an ideal computational structure for determining the more general classes that subsume a given concept within the case base. This structure could easily be used as a measure of similarity between individual and collections of concepts. Further discussion on this benefit can be found in [Wiederhold, 1994].

- *Descriptive*. Unlike a flat database schema, the relationships provided by a semantic information space provide more description of the data field than its name and location. Each concept (and accordingly the name used to describe it) can be given any number of properties that provide both supplemental information as well as the relationship it has to other concepts that may be found in the information space.

- *Longevity*. Formal definitions of information semantics provide a structure that can be represented beyond a specific data source. This allows for concept definitions that extend beyond a current snapshot of the world and enforce a more permanent interpretation. While this does not mean that the meaning of concepts cannot change, it does limit the misunderstandings that can occur based on the use of language alone.

- *Interoperability*. While true context-based interoperability is not yet realized, the structure of semantics provides a more formal basis for promoting predictable data transformation between information spaces. By allowing information spaces to find matching concepts not only by name but also via structural and logical similarities, the likelihood of accurate mapping is increased significantly.

While most of these benefits are tempered by complications they still provide a more reliable basis for reasoning than can be provided by current flat data standards such as RDBMS and XML. This reasoning support provides significant benefit to the DEEP project by enabling query and inference capabilities that are focused on concepts and relations as opposed to words.

However, there are many challenges that must be overcome before semantic technologies can be widely deployed and maintained by an information infrastructure. This includes:

- *Building ontologies*. Ontologies are commonly built as a manual process in which experts and knowledge engineers brainstorm a concept space and begin to formalize definitions. While this proves effective in smaller domains and information spaces it is often the case that the members of larger communities of interest cannot come to an agreement on the definition of some concepts. Depending on the degree of difference between these views it is sometimes impossible to represent both viewpoints in one ontological structure. Some of the varied methodologies for building ontologies are discussed in [Lopez, 1999]

- *Indexing ontologies*. There are currently no standard methods for indexing and allowing searches over ontological concepts and relationships. While many tools allow for a keyword-based search it is often difficult to know exactly what term is used to describe these constructs. Given the ambiguous nature of language there are often many words that could be used to describe one concept. While significant work has been done regarding mapping ontologies to each other [Maedche et al., 2002] little work could be found that looked at an infrastructure to provide access to ontologies that have not yet been directly mapped to the local ontology.

- *Ontology versioning*. Just as the world changes constantly, so should ontologies. While it may hold that cars cannot fly today, someday that may change and an ontology describing cars will have to be re-engineered. However, there may be a situation where a system that has a dependency or reliance on the original ontology to provide the same reasoning constructs it provided before. Therefore, it is vital that a reliable standard for publishing and access of ontology versions

exist to allow for interaction between information spaces without unexpected change. [Klein & Fensel, 2001]

- *Structure vs. Flexibility*. Since ontologies represent an "upper class" of data, it is important to be wary of the commitments made at the ontological level. Any restrictions you place on a concept will be inherited by any instance that is a type of that concept. If an ontology asserts that all cars are blue then it is not possible for an instance of a car under that ontology to be anything but blue. On the other hand if the ontology contains a concept Car that has no properties or relationship it is nothing more than a symbol that provides no logical support. Therefore, there is a balance that must be maintained between the amount of structure you impose at the ontology level and the amount of flexibility given to instances of the concepts that have been described.

During the implementation we found that the landscape of possible specifications was both complex and confusing. For example, it was not initially clear what the differences were between RDF-S terms *subClassOf* and *type*. In addition, RDF-S and OWL share some terms, making it difficult to distinguish between them. Also, we found that the currently available tools are premature and require extensive experience. In fact, RDF generated in one tool will often not open correctly in another tool due to subtle differences between them, even though they are both using the same representation. Further, we believe that RDF encoded in XML format is a misleading representation for human readability even though it is ideal for machine-independent processing. The hierarchical representation reflected by XML encoding is not directly applicable to the structure of the RDF triplets.

However, we believe that the benefits gained should outweigh the shortcomings. Our recommendation to the C2 Community is to invest time and effort into understanding and applying semantic technologies in the form of abstract domain ontologies. More specifically, we suggest using the RDF and RDF-S frameworks as the foundational layer to encode and externalize the semantic concepts within their respective domain. Doing so will enable complex reasoning and portability in the long-term while providing the benefits of lightweight hierarchical reasoning and modularity in the short-term.

4 Conclusion

In October 2007, the DEEP project completed year one of its scheduled four-year effort. The objective of the first year was successfully completed – development of a "research platform" to further support more aggressive research in the areas of distributed C2 and analogical reasoning, and how to apply the technology to advance the state of C2. A key vision of this research in semantic interoperability for distributed planning is to advance current asynchronous "chat room" type collaboration with true distributed planning that can be conducted in parallel.

In this paper we have presented a number of extensions, both existing and planned, to the Object Model Working Group's *CPR* framework. These extensions were made to

support the Distributed Exploratory Episodic Planning (DEEP) project, which is a decision-support planning system designed for providing computer-assisted planning capabilities. Our initial extensions to CPR implemented the ARPI-CPR framework, generating DEEP-CPR, and allowed us to capture planning experiences in both human- and machine-readable form. To support DEEP's experience-based reasoning, much of the knowledge encoded in DEEP-CPR is supported by taxonomies. Our next extensions to DEEP-CPR will incorporate the use of state-of-the-art semantic technologies.

Our extensions to the APRI-CPR model were driven by the need to use a richer content description for reasoning about plans, as well as to be able to encode plans as part of cases in a CBR system. As opposed to the hard-coded planning systems that are currently being used these semantics-based technologies enable more flexible reasoning and expose the structure of stored knowledge in a form that is accessible to both the machine and the users. These capabilities support future USAF C2 requirements for dynamic, distributed heterogeneous planning agents in a more transparent representation that is accessible by all users in the command chain.

References

AF/A5 Plans. May 2006. AF C2 Enabling Concepts.

Alberts, D., and Hayes, R., 2007. Planning: Complex Endeavors. pg. 217, CCRP Publication Series. http://handle.dtic.mil/100.2/ADA465653

Alberts, D., Garstka, J., and Stein, F., 1999. Network-Centric Warfare: Developing and Leveraging Information Superiority. CCRP Publication Series. http://www.dodccrp.org/files/Alberts_NCW.pdf

Alexander, N., Lopez, X., Ravada, S., Stephens, S., and Wang, J., 2004. RDF Data Model in Oracle. Oracle Corporation. http://download-uk.oracle.com/otndocs/tech/semantic_web/pdf/w3d_rdf_data_model.pdf

Berndt, LtGen M. R. (Ret.), 2007. USMC presentation on Commander's Intent (U). Located online in the class section of MAGTF Staff Training Program. http://www.mstp.quantico.usmc.mil/

Brickley, D., and Guha, R., 22 February 2002. Resource Description Framework (RDF) Model and Syntax Specification. W3C Recommendation. http://www.w3.org/TR/1999/REC-rdf-syntax-19990222/

Brickley, D., and Guha, R., 23 January 2003. RDF Vocabulary Description Language 1.0: RDF Schema. W3C Working Draft. http://www.w3.org/TR/2003/WD-rdf-schema-20030123/

Cai, M., and Frank, M., 2004. RDFPeers: A Scalable Distributed RDF Repository based on a Structured Peer-to-Peer Network. Proceedings of *13th International World Wide Web Conference (WWW2004)*, New York.

Klein, M., and Fensel, D., 2001. Ontology Versioning for Semantic Web. Proceedings of *13th Intl' Semantic Web Working Workshop (SWWS'01),* Stanford (July/August 2002), CA, pp. 75–91.

Lopez, F., 1999. Overview of Methodologies for building ontologies. Proceedings of *International Joint Conference on Artificial Intelligence (IJCAI-99)* workshop KRR5, Sweden.

Maedche, A., Motik, B., Silva, N. and Volz, R., 2002. MAFRA – A MApping FRamework for Distributed Ontologies. Proceedings of *13th European Conference on Knowledge Engineering and Knowledge Management EKAW-2002, Madrid, Spain*.

Nejdl, W., Siberski, W., and Sintek., 2003. Design Issues and Challenges for RDF and schema-based peer-to-peer systems. *ACM SIGMOD Record* 32(3): 41-46.

Pan, J., and Horrocks, I., 2004. Owl-e: Extending owl with expressive datatype expressions. Technical report, IMG/2004/KR-SW-01/v1.0, Victoria University of Manchester.

R. Adam Pease, 1998. Core Plan Representation, Object Model Working Group. *Request for comment.*

Pease, A., and Carrico, T., 1997. The JTF ATD Core Plan Representation: A Progress Report. *In proceedings of the AAAI Spring Symposium on Ontological Engineering*.

Pease, R. A., Niles, I., and Li, J., 2002. The suggested upper merged ontology: A large ontology for the semantic web and its applications. In: AAAI-2002 Workshop on Ontologies and the Semantic Web, Working Notes.

US Army. 1997. FM101-5 Staff Organization and Operations Field Manual (1997). Headquarters, Department of Army. http://www.aiai.ed.ac.uk/~arpi/SUO/DOC/fm101-5.pdf

Wiederhold, G., 1994. Interoperation, Mediation and Ontologies. Proceedings of *International Symposium on Fifth Generation Computer Systems (FGCS94),* Tokyo, Japan.